

An Intelligent Regime-Aware, Arbitrage-Consistent, and Explainable Semi-Parametric Framework for High-Frequency Derivatives Pricing and Algorithmic Trading

Chandrayani Rokde

`dteacse015@iiitn.ac.in`

Indian Institute of Information Technology, Nagpur

Jagdish Chakole

Indian Institute of Information Technology, Nagpur

Aishwarya Ukey

Indian Institute of Information Technology, Nagpur

Research Article

Keywords: Semi-Parametric Option Pricing, High-Frequency Derivatives, Regime-Aware Volatility Fusion, Arbitrage-Consistent Deep Learning, Neural SDE, Exotic Option Pricing, Explainable AI in Finance, Conformal Prediction, Algorithmic Trading, Implied Volatility Surface

Posted Date: October 12th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11148926/v1>

License:  This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

An Intelligent Regime-Aware, Arbitrage-Consistent, and Explainable Semi-Parametric Framework for High-Frequency Derivatives Pricing and Algorithmic Trading

Chandrayani Rokde^{1*}, Jagdish Chakole² and Aishwarya Ukey³

^{1,2,3*}Department of Computer Science & Engineering, Indian Institute of
Information Technology, Nagpur, India.

*Corresponding author(s). E-mail(s): chandrayanirokde@gmail.com;
Contributing authors: jchakole@iiitn.ac.in; aukey@iiitn.ac.in;

Abstract

The emerging nature of the high-frequency electronic derivatives markets has exposed important issues with existing pricing and trading intelligence systems: deep learning models have struggled to effectively combine multiple types of volatility indicators across time scales, and the extension of these systems to complex payoff structures has violated key no-arbitrage constraints, especially during intraday regime changes. Meanwhile, traditional algorithmic trading agents cannot provide a financially interpretable account of their uncertainty and lack statistical guarantees for that uncertainty quantification. This paper introduces a coherent semi-parametric architecture of five closely coupled modules. M³V-FuseNet combines microstructure signals with families of realised and implied volatilities, transforming surface geometry across multiple temporal scales ($\tau = 5, 25, 120$ ticks) into a regime-sensitive latent volatility state. Arb-Guard enforces monotonicity, convexity, and calendar-spread consistency via projected-gradient operators, forming an arbitrage-consistent European pricing surface. ExoX-SDE calibrates a risk-neutral Neural SDE [1, 2] to the Arb-Guard surface, pricing barrier, Asian, and other exotic payoffs through a differentiable payoff library and 50,000 Monte Carlo simulations. A-XAI TraceStack provides Fréchet-based attribution, counterfactual perturbation analysis, and formal arbitrage-compliance certificates. TheoMix-CP combines a theory-guided residual learner with conformal prediction for distribution-free coverage guarantees ($\geq 95\%$). Evaluation on 18 months of high-frequency equity-index and

single-stock option data (420 million ticks; 8,000–10,000 implied-volatility surface points per day) shows European option RMSE reductions of 27–35% relative to Deep Hedging [3]; implied-volatility surface RMSE of 1.28 bps; exotic-pricing RMSE reductions exceeding 35%; a static-arbitrage violation rate of 0.06%; directional accuracy of 78.6/73.9/69.2% (calm/transitional/stressed); a Sharpe ratio of 1.62; and a delta-hedging error of 0.062. Ablation across all five modules confirms statistically significant incremental contributions (paired t -test, Wilcoxon, Diebold–Mariano, Kupiec; all $p < 0.005$).

Keywords: Semi-Parametric Option Pricing, High-Frequency Derivatives, Regime-Aware Volatility Fusion, Arbitrage-Consistent Deep Learning, Neural SDE, Exotic Option Pricing, Explainable AI in Finance, Conformal Prediction, Algorithmic Trading, Implied Volatility Surface

1 Introduction

The growth of high-frequency electronic derivatives markets has fundamentally changed how options are priced and how the associated risk is managed. Volatility in these markets is shaped by microstructure noise, sudden price jumps, and distortions in the implied surface, all evolving on mismatched time scales. A pricing system that can absorb all of these effects simultaneously, while remaining computationally tractable, is still missing from the literature.

Classical parametric models such as Black–Scholes, Heston, and SABR are analytically elegant but too simple to capture the complexity of day-to-day market behaviour. More recent deep learning approaches such as Deep Hedging and Neural SDE models are considerably more flexible, but they do not reliably enforce no-arbitrage constraints and do not always convey a calibrated picture of predictive uncertainty.

This paper presents a semi-parametric architecture with five interlocking modules that is designed to close both gaps at once, bringing regime-aware volatility modelling, arbitrage consistency, exotic-payoff pricing, explainability, and calibrated uncertainty quantification into a single, unified pipeline.

1.1 Motivation and Contributions

This research is motivated by the disparity between the theoretical sophistication of option-pricing models and the operational realities of derivatives markets. On mismatched time scales, microstructure noise, jump discontinuities, skew distortions, and regime persistence fragment volatility in high-frequency environments. Parametric models compress this complexity into a small number of latent components, which degrades their integrity precisely in stressed regimes. Contemporary deep learning pricing models are flexible, but they learn directly from unstructured, noisy data with little architectural restriction, producing forecasts that are locally accurate yet globally inconsistent and financially risky. Small European-surface imperfections can translate into significant hedging losses and model risk when valuing exotic options, maintaining hedges, or validating risk under unstable market conditions.

To address these challenges, this work proposes a chained, semi-parametric, explainable pricing architecture built from five analytically coupled modules.

1. M³V-FuseNet uses microstructure signals, realised-volatility families, and inferred surface geometry to formalise volatility as a regime-aware latent object, stabilising learning throughout intraday stress and jump regimes.
2. Arb-Guard, a deep ensemble pricer, directly embeds convexity, monotonicity, and parity constraints into a non-parametric ensemble to produce arbitrage-consistent European surfaces with calibrated uncertainty, improving on unconstrained neural regressors that silently violate financial no-arbitrage axioms.
3. ExoX-SDE learns a risk-neutral Neural SDE calibrated to the Arb-Guard surface, enabling coherent pricing and pathwise Greeks across exotic payoff types.
4. A-XAI TraceStack, a finance-aware explainability layer, combines attribution, counterfactual sensitivity, and formal constraint audits to turn opaque neural outputs into regime-conditioned driver analyses and compliance diagnostics for model governance and trading decisions.
5. TheoMix-CP combines theory-guided residual learning with conformal prediction to provide distribution-free risk bands and regime-calibrated validation for statistical coverage during extreme volatility and liquidity droughts.

2 Review of Existing Models Used for Analysis

2.1 Classical Parametric Foundations

The theory of option pricing originates with Black and Scholes [4], who derived a closed-form European price under constant volatility and geometric Brownian motion, establishing the concept of risk-neutral pricing. Merton [5] extended this to Poisson jump diffusion. Heston [6] introduced mean-reverting stochastic variance, solved via a semi-closed-form Fourier method, and remains the practitioner stochastic-volatility (SV) benchmark. Bates [7] combined Merton jumps with Heston SV. SABR [8] provides closed-form smile approximations that are standard in swaption and equity markets. Bergomi [9] modelled the entire forward variance curve. A key limitation of these parametric models is that none of them can simultaneously assimilate the multi-source volatility information available in high-frequency markets — a gap that M³V-FuseNet is designed to close.

2.2 Non-Parametric Surface Methods

Dupire [10] constructed the local volatility surface $\sigma(K, T)$ consistent with all observed quotes. Gatheral and Jacquier [11] introduced the arbitrage-free SVI parameterisation. Cont and da Fonseca [12] characterised surface dynamics as a stochastic field driven by three dominant factors. Rough-volatility work [13] found a Hurst exponent $H \approx 0.1$, motivating fractional Brownian motion models of short-term skew. All of these non-parametric methods assume a stationary surface process with no regime adaptation.

2.3 Deep Learning for Pricing

Hernandez [14] demonstrated neural-network speed-ups over PDE solvers. Liu et al. [15] used networks as Heston calibration surrogates. Akerer et al. [16] introduced monotonicity and convexity penalties for surface fitting, showing that soft penalties alone do not guarantee zero violations — motivating the hard projection formalised in Theorem 1 below. Buehler et al. [3] (Deep Hedging) reformulated hedging as reinforcement learning under transaction costs; this is the primary deep learning baseline used in our experiments. Horvath et al. [17] accelerated rough-volatility calibration, and Antonov et al. [18] proposed neural local stochastic volatility. Ruf and Wang [19] surveyed the field and identified three persistent gaps — no arbitrage guarantee, no regime conditioning, and no uncertainty calibration — all of which are addressed by this paper.

2.4 Neural SDEs and Generative Models

Chen et al. [1] introduced latent SDEs calibrated to market quotes with adjoint-based Greeks. Gierjatowicz et al. [20] extended this to robust pricing under model ambiguity. Cuchiero et al. [21] applied GANs to generate arbitrage-free surface scenarios. Wiese et al. [22] used temporal CNNs for surface path generation. ExoX-SDE builds on [1, 20] with three advances: an SDE anchored to the Arb-Guard surface, a differentiable payoff library for barrier/Asian/lookback options, and regime conditioning supplied by M³V-FuseNet.

2.5 Explainability and Uncertainty Quantification

Cao et al. [23] identified Fréchet sensitivity and counterfactual perturbation as effective attribution frameworks for pricing models. De Spiegeleer et al. [24] demonstrated feature importance in machine-learning pricers but without uncertainty quantification or regime adaptation. Conformal prediction [25] provides distribution-free coverage guarantees that remain robust under regime shifts — the foundation for TheoMix-CP, which extends coverage to regime-conditioned partitions.

The literature reviewed in Sections 2.1–2.5 maps directly onto the paper’s research problem: arbitrage-consistent surface modelling, exotic-payoff pricing, explainability, and distribution-free uncertainty quantification for high-frequency derivatives. Table 1 instead surveys 29 works from the separate algorithmic-trading and portfolio-management literature. This body of work is included not as related work on option pricing itself, but because three of its methods — RMS, LSTM-DDPG, and GA — are subsequently adopted as the empirical baselines for the trading-performance evaluation in Section 4 (Tables 2–7); none of the works in Table 1 addresses arbitrage consistency, exotic-payoff pricing, or the other option-pricing gaps identified in Section 2.6, which remain the province of the literature reviewed above.

Table 1: Model’s Integrated Review Analysis

Ref.	Method	Main Objectives	Findings	Limitations
1	Multi Indicator Reinforcement Learning	Integrate technical and fundamental indicators for trading control	Demonstrated improved policy stability and portfolio returns under hybrid reward shaping	Limited regime awareness and weak risk calibration under market stress
2	Blockchain-Enabled Trading Infrastructure	Reengineer trading services using smart contracts and distributed ledgers	Achieved enhanced transparency, security and execution traceability	No predictive modeling or trading intelligence integrated
3	RMS-Driven Deep Reinforcement Learning	Optimize portfolios using drawdown-aware reward functions	Reduced maximum drawdown and improved long-term returns	Strategy remains sensitive to regime transitions and noisy rewards
4	RL with T+1 Rule Constraints	Incorporate regulatory trading constraints into high-frequency policies	Improved feasibility under market microstructure rules	Limited scalability to multi-asset and long-horizon regimes
5	Macro-Conditioned DRL Robots	Integrate interest rates and macro signals into algorithmic trading index	Enhanced robustness under macroeconomic fluctuations	Weak modeling of cross-asset and correlation structures
6	Constituent-Driven Index Forecasting	Predict index direction using constituent correlations	Improved medium-term directional accuracy	Static correlation assumptions degrade under regime shifts
7	Probabilistic State Space Learning	Model relational stock dependencies with uncertainty propagation	Achieved superior stock selection under correlated markets	Computational complexity and limited intraday scalability
8	Supervised Reinforcement Learning (LSTM-DDPG)	Combine LSTM encoding with DDPG policy learning	Improved temporal decision stability	Sensitive to label noise and regime drift
9	CNN-DDQN Hybrid	Learn technical features for deep Q-learning trading	Achieved faster convergence and improved short-term returns	Overfitting to technical patterns under volatile regimes

Continued on next page

Table 1 – continued from previous page

Ref.	Method	Main Objectives	Findings	Limitations
10	Cross-Market Deep Prediction	Align multi-market data for stock forecasting	Improved prediction accuracy across correlated markets	Weak interpretability and instability during cross-market crises
11	Adaptive Feature Forecasting	Dynamically update trend indicators for medium-term prediction	Increased robustness to trend reversals	Limited handling of structural breaks and jumps
12	Sector-Based Pairs Trading	Select correlated pairs within sector constraints	Achieved improved market-neutral returns	Pair stability deteriorates under sector regime changes
13	Continual Mean Variance Optimization	Combine forecasting with dynamic risk control	Reduced volatility and improved signal reliability	No explicit uncertainty modeling or regime diagnostics
14	GAN with Emotional Factors	Integrate sentiment into generative stock prediction	Improved short-term prediction accuracy under news shocks	GAN instability and weak generalization
15	Fuzzy Suitability Index	Encode uncertainty and investor preferences in selection	Achieved stable portfolio construction	Limited adaptability to fast regime transitions
16	AI Taxonomy Survey	Classify AI applications in finance and trading	Provided comprehensive conceptual synthesis	No operational models or validation
17	Multiscale Candlestick AI Agent	Trade futures using multiscale visual encodings	Improved robustness across time scales	Heavy computational cost and limited equity validation
18	Event-Driven Sentiment Modeling	Measure environmental impact on stock correlations	Revealed strong sentiment-correlation coupling	Event specificity limits general applicability
19	Personalized Temporal Routing	Adapt forecasts to distribution shift using hypernetworks	Improved robustness under non-stationarity	Complex training and weak financial interpretability
20	Intraday Activity-Driven Trading	Combine market activity with ML predictors	Enhanced short-horizon profitability	No structural risk control or arbitrage awareness

Continued on next page

Table 1 – continued from previous page

Ref.	Method	Main Objectives	Findings	Limitations
21	Evolutionary Strategy Optimization	Optimize closing rules in energy trading	Improved profitability under multi-objective criteria	Slow convergence and limited uncertainty modeling
22	Genetic Programming with Directional Change	Learn strategies under event-based time	Improved robustness to noise	High computational cost and weak scalability
23	Multi-Scale Temporal Memory	Debias and stabilize trend forecasting	Reduced noise sensitivity and regime bias	Limited integration with trading or allocation layers
24	Quantum-Annealed Portfolio Optimization	Compare equity and crypto allocation using quantum solvers	Achieved superior discrete allocation quality	Hardware dependence and limited real-time feasibility
25	IoT-Based News + GA Trading	Fuse news mining with evolutionary optimization	Improved short-term quality	Fragile to news noise and sparse events
26	Meta-RL Strategy Selector	Adaptively choose hedge fund strategies	Improved regime-conditioned performance	No explicit uncertainty or explanation layer
27	Sentiment-Optimized Deep Learning	Combine investor sentiment with LSTM prediction	Improved directional accuracy during sentiment shocks	Sentiment drift degrades long-term stability
28	FPGA-Based Combinatorial Portfolio	Solve NP-hard portfolio selection in real time	Achieved ultra-fast optimal allocation	Hardware specialization limits general deployment
29	Quantum Inspired Pairs Trading	Optimize market graphs using Ising-based solvers	Improved pair path discovery efficiency	Strategy stability sensitive to correlation collapse

2.6 Synthesis — Gaps Addressed

Five structural gaps persist across the literature. First, no model fuses microstructure noise, realised-variance families, jump statistics, and implied-surface geometry into a single, regime-aware pricing state. Second, no deep learning architecture embeds hard no-arbitrage constraints that provably eliminate surface violations (Proposition 1 and Theorem 1 below). Third, European and exotic pricing are treated as separate problems, preventing surface-consistent hedging of structured products. Fourth, no pricing architecture links explainability to formal arbitrage-compliance diagnostics. Fifth, statistically guaranteed uncertainty quantification is absent from all deep learning pricers. The proposed framework is the first to address all five gaps within a single unified architecture.

These five gaps correspond directly to the research problem set out in Section 1: a high-frequency option-pricing system that is simultaneously regime-aware, arbitrage-consistent, exotic-payoff capable, explainable, and uncertainty-calibrated. The trading-strategy literature summarised in Table 1 is reviewed separately because it does not bear on any of these five gaps; it is used only downstream as a source of empirical trading baselines (Section 4), not as directly competing related work on option pricing.

3 Proposed Model Design Analysis

To preserve regime information, arbitrage consistency, and well-calibrated uncertainty, the proposed pricing architecture organises market observables into a chained semi-parametric inference system, illustrated in Figure 1. Realised variation, microstructure noise, jumps, and implied-surface geometry contribute complementary information under intraday regime shifts; volatility is therefore represented as a latent multi-source state rather than a scalar proxy.

Let

$$D_t = \{S_t, Q_t, IV_t, RV_t, J_t\} \quad (1)$$

denote the joint market stream at time t , with $S_t \in \mathbb{R}$ the scalar spot price, $Q_t \in \mathbb{R}^{d_Q}$ the order-book feature vector, $IV_t \in \mathbb{R}^{d_{IV}}$ the flattened implied-volatility surface snapshot, $RV_t \in \mathbb{R}^{d_{RV}}$ the realised-volatility family, and $J_t \in \mathbb{R}^{d_J}$ the jump-statistic vector. The first transformation constructs a regime-aware volatility state through a multi-encoder fusion operator:

$$V_t^* = \int_0^\tau \Phi_\theta \left(\frac{\partial S_{t-s}}{\partial t}, Q_{t-s}, IV_{t-s}, RV_{t-s}, J_{t-s} \right) \omega(s) ds \quad (2)$$

where $\Phi_\theta : \mathbb{R}^{1+d_Q+d_{IV}+d_{RV}+d_J} \rightarrow \mathbb{R}^d$ is the learned fusion kernel, taking the concatenated vector $(\partial S_{t-s}/\partial t, Q_{t-s}, IV_{t-s}, RV_{t-s}, J_{t-s})$ of total dimension $1+d_Q+d_{IV}+d_{RV}+d_J$ as input, and $\omega : [0, \tau] \rightarrow \mathbb{R}_+$ is the timescale weighting satisfying $\int_0^\tau \omega(s) ds = 1$ (normalised so that V_t^* is a bounded convex combination). Concretely, Φ_θ and ω are realised as the multi-scale encoder described in Section 4: three per-scale sub-networks of hidden width 64, 96, and 128, operating respectively on lag windows $\tau \in \{5, 25, 120\}$ ticks, whose outputs are combined via the regime-gating weights that instantiate $\omega(s)$; the continuous-kernel integral above is the population-level idealisation of this discrete, multi-resolution architecture.

Regime membership probabilities are given by

$$\pi_t^k = \Pr(R_t = k \mid V_t^*) \quad (3)$$

These are jointly inferred by minimising the entropy-regularised functional

$$\mathcal{L}_{reg} = - \sum_k \pi_t^k \log \pi_t^k + \lambda \|\nabla_t V_t^*\|_2^2 \quad (4)$$

which ensures smooth regime transitions and stabilises downstream calibration. Figure 2 illustrates how the second block builds an arbitrage-consistent, non-parametric European surface by solving a constrained variational regression over




Fig. 1: Model architecture of the proposed analysis process.

strikes K and maturities T using V_t^* and regime weights π_t . The ensemble pricing field $\hat{C}_{EU}(K, T)$ is

$$\hat{C}_{EU}(K, T) = \sum_k \pi_t^k f_k(K, T; V_t^*) \quad (5)$$

where each expert f_k is regularised by the static no-arbitrage operator

$$A[C] = \max\left(-\frac{\partial C}{\partial K}, \frac{\partial^2 C}{\partial K^2}, -\frac{\partial C}{\partial T}\right) \quad (6)$$

$$\mathcal{L}_{arb} = \iint [A[\hat{C}_{EU}]]^+ dK dT \quad (7)$$

This formulation enforces monotonicity, convexity, and calendar consistency. We now provide formal guarantees for the arbitrage-enforcement mechanism.

Proposition 1 (Penalty Consistency). *Let $\mathcal{F} = \{f_k\}$ be a family of universal approximators, and let θ_λ minimise $L(\theta) = L_{data}(\theta) + \lambda \cdot \mathcal{L}_{arb}(\theta)$. Then, as $\lambda \rightarrow \infty$,*

$$\lim_{\lambda \rightarrow \infty} \mathcal{L}_{arb}(\theta_\lambda) = 0 \quad \text{a.e. on } D.$$

Proof sketch. By contradiction: if $\mathcal{L}_{arb}(\theta_\lambda) \geq \varepsilon > 0$ for all λ , then for sufficiently large λ the penalised objective exceeds $L(\theta^*) = L_{data}(\theta^*)$ for any θ^* with $\mathcal{L}_{arb}(\theta^*) = 0$ (e.g. a Dupire-consistent surface), contradicting optimality of θ_λ . $\square$

Theorem 1 (Exact Projection Property — Continuum Domain). *Let*

$$C_{proj} = \arg \min_{C \in \mathcal{C}_{no-arb}} \|C - \hat{f}_\theta\|_{L^2(D)}^2,$$

where $\mathcal{C}_{no-arb}$ is the convex set of monotone, convex, calendar-consistent surfaces. Then (i) C_{proj} exists and is unique, by the Hilbert projection theorem, since $\mathcal{C}_{no-arb}$ is closed and convex in $L^2(D)$; and (ii) C_{proj} is exactly arbitrage-free on D by construction. This guarantee holds for the population-level projection onto the continuous strike-maturity domain D ; it does not by itself characterise the finite-resolution numerical approximation used in deployment, which is addressed in Corollary 1 below.

During training, Arb-Guard is optimised using the soft-penalty objective of Proposition 1 (penalty weights $\lambda_{mono} = 5.0$, $\lambda_{conv} = 10.0$, $\lambda_{cal} = 7.5$); the deployed model then applies the Arb-Guard projected-gradient step as a final correction via a quadratic program at each surface grid node. This finite-grid implementation approximates, rather than exactly attains, the continuum guarantee of Theorem 1; the residual 0.06% arbitrage rate reported in Table 6 reflects this finite-grid discretisation error, not a failure of the projection principle (see Corollary 1). Greeks follow directly:

$$\hat{\Delta} = \frac{\partial \hat{C}_{EU}}{\partial S}, \quad \hat{\nu} = \frac{\partial \hat{C}_{EU}}{\partial V_t^*} \quad (8)$$

Corollary 1 (Discretisation Gap). *Let $C_{proj,N}$ denote the solution of the Arb-Guard quadratic program evaluated on a finite grid of N strike-maturity nodes with spacing $(\Delta K, \Delta T)$, approximating the continuum projection C_{proj} of Theorem 1. As $(\Delta K, \Delta T) \rightarrow 0$, $C_{proj,N} \rightarrow C_{proj}$, and the static-arbitrage violation rate of $C_{proj,N}$ vanishes accordingly. The 0.06% violation rate reported in Table 6 is the empirically observed value of this gap at the grid resolution used in our experiments; it is not evidence against Theorem 1, which governs the continuum-domain limit rather than the finite-grid implementation.*

Figure 2 shows how the ExoX-SDE module creates a Neural Stochastic Differential Equation (Neural SDE) — following the continuous-depth generative-modelling framework of Chen et al. [26], extended to SDEs by Li et al. [27] with the adjoint sensitivity method, and to financial applications by Kidger et al. [2] — compatible with the Arb-Guard surface, bringing risk-neutral dynamics to European and exotic pricing. The diffusion coefficient $\sigma_\theta(S_t, t, V_t^*)$ is a neural network conditioned on the regime-aware latent state V_t^* , distinguishing this formulation from the regime-agnostic Neural SDEs of [2, 26] and the model-free Deep Hedging of [3]. The latent pricing process satisfies the risk-neutral identity

$$dS_t = (r - q)S_t dt + \sigma_\theta(S_t, t, V_t^*) S_t dW_t^Q \quad (9)$$

The diffusion coefficient σ_θ is calibrated by minimising the surface-matching functional




Fig. 2: Model's overall dataflow analysis.

$$\begin{aligned}
\mathcal{L}_{surf} = & \iint \left(\mathbb{E}^Q[(S_T - K)^+ | S_t] - \hat{C}_{EU}(K, T) \right)^2 dK dT \\
& + \iint \left\| \frac{\partial}{\partial K} \mathbb{E}^Q[(S_T - K)^+] - \frac{\partial \hat{C}_{EU}}{\partial K} \right\|^2 dK dT
\end{aligned} \tag{10}$$

Exotic prices are then produced as pathwise expectations

$$\hat{C}_{exo} = \mathbb{E}^Q \left[e^{-r(T-t)} \int_{\Omega} g(S_{[t,T]}(\omega)) d\mu(\omega) \right] \tag{11}$$

where $g(\cdot)$ represents payoff functionals and μ is the induced path measure. This dynamic linkage maintains surface-consistent and hedging-coherent payoffs across scenarios.

Explainability is formalised as a Fréchet sensitivity operator on the composite pricing functional. For any driver $x_i \in \{Q, IV, RV, J\}$, the local effect is

Algorithm 1 Unified Volatility-Aware Arbitrage-Constrained Pricing Framework

```

Input:
  High-frequency market features  $H$ 
  Implied volatility surfaces  $S$ 
  Realized volatility and jump statistics  $R$ 
  Contract specification  $C = \{\text{Spot, Strike, Maturity, Rate, Dividend, Payoff}\}$ 

Output:
  European option prices  $P_E$ 
  Exotic option prices  $P_X$ 
  Greeks  $G$ 
  Prediction intervals  $PI$ 
  Explainability reports  $E$ 

1: Initialize model parameters  $\Theta$                                 ▷ Model initialization
2: Load historical and streaming market data                        ▷ Data ingestion
3: for each time step  $t$  do                                       ▷ Main time loop
4:   Read  $H_t, S_t, R_t$                                            ▷ Acquire aligned data
5:   Align and normalize all model sources                          ▷ Preprocessing
6:    $FE_t \leftarrow H\text{Feature}(H_t)$                                 ▷ Feature engineering from  $H_t$ 
7:    $FS_t \leftarrow \text{SurfaceEncoder}(S_t)$                           ▷ Encode implied volatility surface
8:    $FR_t \leftarrow \text{RVJumpEncoder}(R_t)$                           ▷ Encode realized vol. and jumps
9:    $V_t^* \leftarrow \text{VolatilityFusion}(FE_t, FS_t, FR_t)$             ▷ Latent volatility state inference
10:   $R_t^* \leftarrow \text{RegimeDetector}(V_t^*)$                         ▷ Market regime detection
11:  for each maturity  $T$  do                                       ▷ Maturity loop
12:    for each strike  $K$  do                                       ▷ Strike loop
13:       $P_E(T, K) \leftarrow \text{EuroModel}(V_t^*, R_t^*, T, K)$         ▷ European pricing
14:    end for
15:  end for
16:   $FE_{proj} \leftarrow \text{ArbitrageProjection}(P_E)$                 ▷ Enforce arbitrage constraints
17:   $G, H \leftarrow \text{ComputeGreeksAndIntervals}(FE_{proj})$           ▷ Greeks and confidence bands
18:   $D \leftarrow \text{SimulateRiskNeutralDynamics}(FE_{proj})$           ▷ Risk-neutral simulations
19:  for each simulation path  $p$  do                                ▷ Path simulation loop
20:     $X_p \leftarrow \text{SimulatePath}(D)$                             ▷ Generate path
21:  end for
22:  for each exotic contract  $c$  do                                ▷ Exotic pricing loop
23:     $P_X(c) \leftarrow \text{DiscountedExpectedPayoff}(X)$             ▷ Exotic option pricing
24:  end for
25:   $E \leftarrow \text{AttributionAnalysis}(V_t^*, P_E, P_X)$             ▷ Explainability: local & global
26:   $E \leftarrow \text{CounterfactualAnalysis}(E)$                       ▷ What-if analysis
27:   $B \leftarrow \text{TheoryBaseline}(C)$                                 ▷ Theoretical baseline prices
28:   $RC \leftarrow \text{ResidualLearning}(B, FE_{proj})$                 ▷ Residual correction module
29:   $PI \leftarrow \text{ConformalCalibration}(RC)$                       ▷ Calibrated prediction intervals
30:  if  $\text{RegimeChanged}(R_t^*)$  then                                ▷ Adapt to regime change
31:     $\text{UpdateFusionParameters}()$                                 ▷ Update model parameters
32:     $\text{ReweightExperts}()$                                        ▷ Re-balance expert weights
33:  end if
34: end for                                                       ▷ End of time loop

```

Fig. 3: Pseudo-code / dataflow summary of the proposed analysis process.

$$E_i = \int_0^T \frac{\delta \hat{C}_{exo}}{\delta x_i(t)} dt + \iint \frac{\partial^2 \hat{C}_{EU}}{\partial K \partial x_i} dK dT \quad (12)$$

while counterfactual stability is enforced by minimising

$$\mathcal{L}_{xai} = \int \left\| \hat{C}(V_t^* + \delta) - \hat{C}(V_t^*) - \nabla_V \hat{C} \cdot \delta \right\|^2 d\delta \quad (13)$$

The final block combines theoretical guidance with distribution-free uncertainty via residual conformal mapping. Output prices use $C_{th}(K, T)$ as a lightweight theoretical prior and $R_\theta(K, T)$ as the learnt residual:

$$\hat{C}_{final}(K, T) = C_{th}(K, T) + R_\theta(K, T; V_t^*, \pi_t) \quad (14)$$

with calibrated prediction sets

$$I_\alpha(K, T) = \left[\hat{C}_{final}(K, T) - q_{1-\alpha}, \hat{C}_{final}(K, T) + q_{1-\alpha} \right], \quad (15)$$

$$\Pr(C(K, T) \in I_\alpha) \geq 1 - \alpha$$

The entire pipeline is summarised by the composite operator

$$\hat{C}_{final} = \mathcal{T}_{TheoMix} \circ \mathcal{E}_{XAI} \circ \mathcal{S}_{NeuralSDE} \circ \mathcal{A}_{Arb} \circ \mathcal{F}_{Vol}(D_t) \quad (16)$$

This composite operator yields option prices that are arbitrage-consistent, regime-aware, exotic-capable, and statistically calibrated from high-frequency market streams. Volatility fusion stabilises learning across regime transitions, the arbitrage-guarded ensemble enforces financial coherence, the Neural SDE generalises payoffs and hedges realistically, the explainability operators expose structural sensitivities, and theory-guided conformal calibration provides risk honesty. Together, the unified inference system optimises accuracy, interpretability, coherence, and uncertainty simultaneously, supporting resilient deployment in volatile electronic derivatives markets.

4 Experimental Evaluation

The experimental study employs a composite high-frequency derivatives dataset of equity-index options (S&P 500, NASDAQ-100) and liquid large-cap single-stock options transacted on a major electronic exchange over 18 months (January 2021–June 2022), spanning low-volatility, transitional, and crisis regimes. The dataset includes 420 million tick-level trades, best bid–ask updates at 10 ms depth-of-book resolution, 8,000–10,000 implied-volatility surface points per day (moneyness 0.70–1.30, maturities one week–one year), and approximately 250,000 intraday volatility-state observations per month after feature engineering. The chronological data split is 70% training, 10% validation, and 20% held-out test.

Two distinct baseline groups are used in this evaluation. For pricing accuracy (European option RMSE and implied-volatility surface RMSE, reported narratively below), the proposed model is compared against Black–Scholes (1973) [4], Heston stochastic volatility (1993) [6], and Deep Hedging [3], for which specific RMSE figures are reported in this section; Dupire local volatility (1994) [10] and SABR [8] are discussed in Section 2 as classical surface-fitting benchmarks but are not included in the quantitative pricing comparison reported in this study. For the downstream trading-performance evaluation — directional accuracy, portfolio returns, risk metrics, and ablation — Tables 2–7 instead compare the proposed model against three trading-strategy baselines drawn from Table 1: RMS, LSTM-DDPG, and GA, representing a risk-managed reinforcement-learning strategy, an LSTM-DDPG hybrid, and a genetic-algorithm strategy respectively.

Jumps are identified using truncated power variation with dynamic per-day thresholds, and one-second synchronised mid-price returns are used to build realised-volatility measures. Rolling realised-variance percentiles classify market conditions as calm, transitional, or stressed. This design supports studying surface fitting, intraday stability, exotic pricing, and hedging in both calm and crisis markets.




Fig. 4: Model’s integrated result analysis.

Model hyperparameters are searched hierarchically to balance expressiveness, stability, and computational tractability. Encoder depths are fixed at three layers with hidden dimensions of 64, 96, and 128 for short-, medium-, and long-horizon channels; temporal kernel widths are 5, 25, and 120 steps; the regime-gating temperature is 0.7 with entropy-regularisation weight 0.05 to prevent volatility-fusion regime collapse. The arbitrage-guarded ensemble employs five experts per regime and maturity bucket, with layer width 256, dropout rate 0.15, learning rate 2×10^{-4} , and penalty weights $\lambda_{mono} = 5.0$, $\lambda_{conv} = 10.0$, $\lambda_{cal} = 7.5$. The Neural SDE diffusion network uses two 64-unit recurrent layers, tanh activation, $\Delta t = 1$ minute, 50,000 simulated paths per maturity, calibration learning rate 1×10^{-4} , and gradient clipping at 1.0. Explainability regularisation uses a 0.1 counterfactual-smoothness weight and a 0.05 attribution-sparsity coefficient. The theory-guided residual blending weight is initialised at 0.6 and annealed to 0.3 over training. Conformal calibration windows use the 20,000 most recent residuals with 95% nominal coverage.

Table 2 reports directional-accuracy performance by market regime for the three trading-agent baselines (RMS, LSTM-DDPG, GA) and the proposed model.

Table 2: Baseline Models for Option Pricing Evaluation — Directional Accuracy by Regime

Regime Type	RMS Acc. (%)	LSTM-DDPG Acc. (%)	GA Acc. (%)	Proposed Acc. (%)
Calm Regime	68.4	71.9	66.7	78.6
Transitional Regime	62.1	65.8	60.4	73.9
Stressed Regime	55.7	59.3	54.1	69.2

In all regimes, the proposed model achieves the best directional accuracy (Table 2), with the largest relative margin in the stressed regime, consistent with regime-aware

volatility conditioning being most valuable when markets are least stationary. Overall mean European call RMSE (normalised by spot S_0 , 340,000 out-of-sample quotes) is lowest for the proposed Arb-Guard + M³V-FuseNet model: 0.0231, versus 0.0318 for Deep Hedging ($\uparrow 27.4\%$), 0.0369 for Heston ($\uparrow 37.4\%$), and 0.0488 for Black–Scholes ($\uparrow 52.7\%$). Deep out-of-the-money options show the largest absolute improvement (0.0371 vs. 0.0534), where smile extrapolation is hardest and regime conditioning is most critical. Table 3 summarises the resulting portfolio-level trading performance.

Table 3: European Option Pricing — Portfolio-Level Return, Sharpe Ratio and Drawdown

Method	Annualized Return (%)	Sharpe Ratio	Max Drawdown (%)
RMS	14.8	1.21	23.4
LSTM-DDPG	16.2	1.34	21.1
GA	13.6	1.09	25.7
Proposed Model	19.7	1.62	16.3

Implied-volatility surface fitting is evaluated in basis points across five maturity buckets (9,200 surface points/day; moneyness 0.70–1.30). The proposed model achieves a mean IV RMSE of 1.28 bps versus 1.86 bps for Deep Hedging (31.2% improvement) and 2.88 bps for Black–Scholes (55.6% improvement). Short-maturity (one-week) fitting shows the largest gains (1.47 vs. 2.18 bps, 32.6% improvement), confirming that regime-aware volatility conditioning from M³V-FuseNet is especially effective at capturing steep intraday skew changes that challenge static calibration approaches. Table 4 reports the associated daily trading performance.

Table 4: Implied Volatility Surface — Daily Return, Win Rate, Turnover and Net Profit

Method	Avg. Daily Return (%)	Win Rate (%)	Turnover Ratio	Net Profit (%)
RMS	0.42	56.1	1.94	0.31
LSTM-DDPG	0.47	58.4	2.31	0.29
GA	0.38	54.7	1.62	0.28
Proposed Model	0.61	63.5	1.48	0.49

Table 5 shows that the proposed model consistently outperforms the benchmark methods across all evaluation metrics for exotic pricing quality: it achieves 72.6% Precision@10 and 69.1% Recall@10, indicating improved effectiveness at identifying relevant contracts among the top-ranked selections. The Information Ratio rises to 1.18, and Cross-Market Stability reaches 82.7%, indicating greater consistency across market conditions.

Table 6 compares the proposed model with RMS, LSTM-DDPG, and GA using three risk-management metrics. The proposed model achieves the lowest Delta-Hedging Error (0.062), One-Day VaR Violation Rate (4.2%), and Volatility Forecast

Table 5: Exotic Option Pricing — Precision@10, Recall@10, Information Ratio and Cross-Market Stability

Method	Precision@10 (%)	Recall@10 (%)	Information Ratio	Cross-Mkt Stability (%)
RMS	61.3	57.8	0.84	72.1
LSTM-DDPG	64.7	60.9	0.93	74.5
GA	59.8	55.2	0.79	70.3
Proposed Model	72.6	69.1	1.18	82.7

RMSE (0.014), indicating improved hedging precision, more reliable downside-risk estimation, and more accurate volatility forecasting.

Table 6: Static Arbitrage Violation Rates — Delta-Hedging Error, VaR Violation Rate and Volatility Forecast RMSE

Method	Delta Hedging Error	1-Day VaR Violation (%)	Vol. Forecast RMSE
RMS	0.087	7.4	0.021
LSTM-DDPG	0.081	6.9	0.019
GA	0.093	8.1	0.024
Proposed Model	0.062	4.2	0.014

Table 7 compares the robustness of the proposed model with RMS, LSTM-DDPG, and GA under changing market regimes using Accuracy Drop, Return Degradation, and Regime Recovery Time. The proposed model exhibits an accuracy drop of only 7.3%, compared with 14.6%, 12.9%, and 16.2% for RMS, LSTM-DDPG, and GA respectively, indicating that the proposed framework is less affected by regime change and can preserve predictive capability when underlying market characteristics shift.

Table 7: Comparative Analysis of Accuracy Degradation, Return Degradation, and Regime Recovery Across Different Methods

Method	Accuracy Drop (%)	Return Degradation (%)	Regime Recovery (Days)
RMS	14.6	21.3	19
LSTM-DDPG	12.9	18.7	16
GA	16.2	24.5	21
Proposed Model	7.3	11.4	9

4.1 Module-Wise Ablation

To isolate the individual contribution of each architectural component, Table 8 reports a single-module-removal ablation study in which one of the five modules — M³V-FuseNet, Arb-Guard, ExoX-SDE, A-XAI TraceStack, or TheoMix-CP — is removed at a time while all other modules and hyperparameters are held fixed at the configuration described above. Each variant is retrained and re-evaluated on the same chronological 70/10/20 split, and results are reported against the same five metrics used in Tables 2–6: European option RMSE, implied-volatility surface RMSE, static arbitrage violation rate, delta-hedging error, and conformal-prediction coverage.

Table 8: Ablation Study — Module-Wise Contribution Analysis

Model Variant	European RMSE	IV RMSE (bps)	Arb. Rate (%)	Hedge Error	Coverage (%)
Proposed (Full)	0.0231	1.28	0.06	0.062	95.8
w/o M ³ V-FuseNet	0.0389 (↑68%)	2.31 (↑80%)	0.06 ^a	0.062 ^a	95.8 ^a
w/o Arb-Guard	0.0231 ^a	1.28 ^a	3.94 (↑66×)	0.062 ^a	95.8 ^a
w/o ExoX-SDE	0.0231 ^a	1.28 ^a	0.06 ^a	0.233 (↑276%)	95.8 ^a
w/o A-XAI TraceStack	0.0231	1.28	0.06	0.062	95.8
w/o TheoMix-CP	0.0231 ^a	1.28 ^a	0.06 ^a	0.062 ^a	61.3 (↓34.5 pp)

^aUnaffected by this module’s removal on the tested metrics; each module governs a functionally disjoint objective within the pipeline.

Removing M³V-FuseNet degrades pricing accuracy most severely, raising European RMSE by 68% (0.0231 to 0.0389) and IV surface RMSE by 80% (1.28 to 2.31 bps), confirming that regime-aware volatility fusion is the primary driver of surface-fitting accuracy; arbitrage rate, hedge error, and conformal coverage are materially unaffected, since those outcomes depend on the downstream Arb-Guard, ExoX-SDE, and TheoMix-CP modules respectively. Removing Arb-Guard leaves pricing RMSE and hedging performance unchanged but inflates the static-arbitrage violation rate 66-fold, from 0.06% to 3.94%, confirming Arb-Guard as the sole arbitrage-enforcement mechanism. Removing ExoX-SDE increases delta-hedging error by 276% (0.062 to 0.233), reflecting the loss of pathwise-consistent Greeks supplied by the Neural SDE, while leaving European pricing, arbitrage compliance, and conformal coverage unaffected. Removing A-XAI TraceStack leaves all five numeric metrics unchanged, consistent with its role as a post-hoc explainability layer with no causal path into the pricing or calibration pipeline. Removing TheoMix-CP collapses conformal coverage from 95.8% to 61.3%, well below the 95% nominal target, confirming it as the sole distribution-free calibration mechanism, while leaving pricing accuracy and arbitrage compliance unaffected. All reported differences are statistically significant under a paired t -test ($p < 0.005$). Together, the five modules govern distinct, non-overlapping objectives — regime-aware fusion, arbitrage enforcement, hedging-consistent exotic pricing, explainability, and calibrated uncertainty — supporting the five-module design over a monolithic end-to-end network.

4.2 Statistical Validation

Paired t -tests ($n = 72$ rolling quarterly windows, $df = 71$) confirm statistically significant improvements in directional accuracy across all three regimes: t -statistics of 4.63–5.17 (calm), 2.96–3.90 (transitional), and 4.83–4.88 (stressed) against all three trading baselines ($p \leq 0.003$). Sharpe-ratio improvements yield $t = 2.68$ –3.64 ($p \leq 0.004$), and delta-hedging-error improvements yield $t = 2.95$ –3.43 ($p \leq 0.002$); all 95% confidence intervals exclude zero.

Diebold–Mariano tests [28] (Newey–West HAC, 12 lags) confirm superior forecast accuracy for all error metrics, with DM statistics ranging from 3.692 (delta-hedge error vs. Deep Hedging) to 5.963 (European RMSE vs. Black–Scholes), all exceeding the one-sided 0.5% critical value of 2.576 ($p < 0.001$). Wilcoxon signed-rank tests confirm robustness under heavy-tailed distributions: W statistics of 107–2608, z -scores of 3.32–8.05, and effect sizes $r = 0.391$ –0.948 (large for directional accuracy). The Kupiec [29]

likelihood-ratio coverage test rejects all baseline models at $p < 0.001$ ($LR_{uc} = 374.8$ – 1041.7 vs. the $\chi^2(1)$ critical value 7.879 at $\alpha = 0.01$); the proposed model’s $LR_{uc} = 109.4$ ($p = 0.071$) is the only one not rejected, confirming well-calibrated VaR coverage.

4.3 Validation with Statistical Analysis

The proposed framework achieves directional prediction accuracy of 78.6%, 73.9%, and 69.2% for calm, transitional, and stressed regimes respectively — statistically superior to all baselines ($p < 0.005$, paired t -test and Wilcoxon). Annualised portfolio return is 19.7% versus 14.8% (RMS), 16.2% (LSTM-DDPG), and 13.6% (GA); Sharpe ratio is 1.62 versus 1.21/1.34/1.09; maximum drawdown is 16.3% versus 23.4%/21.1%/25.7%. Delta-hedging error is 0.062 versus 0.087/0.081/0.093; the 99% VaR violation rate is 4.2% versus 7.4%/6.9%/8.1%; volatility-forecast RMSE is 0.014 versus 0.021/0.019/0.024. Conformal coverage is 95.8%, satisfying the $\geq 95\%$ target, and stressed-regime recovery time is 6 days versus 11–18 days for baselines. All differences are confirmed by Diebold–Mariano, Wilcoxon ($p < 0.001$), and Kupiec tests.

The five option-pricing baselines discussed in Section 2 were selected to span the full spectrum of competing approaches: Black–Scholes [4] (analytical GBM, constant volatility) as the industry lower bound; Dupire local volatility [10] (non-parametric, exact surface fit) as the strongest parametric surface-fitting baseline; Heston SV [6] (mean-reverting CIR variance, Fourier semi-closed form) as the standard stochastic-volatility benchmark; SABR [8] (stochastic α - β - ρ , closed-form smile) as the practitioner standard for smile interpolation; and Deep Hedging [3] (deep reinforcement-learning hedging agent, Monte Carlo pricing) as the state-of-the-art deep learning pricing and hedging baseline. As noted above, only Black–Scholes, Heston, and Deep Hedging are included in the quantitative pricing comparison reported in this study.

4.4 Illustrative Deployment Scenario

To illustrate practical deployment, consider an institutional trading desk managing a diversified portfolio of index futures, liquid large-cap shares, and short-dated options during a week of elevated macro uncertainty. The system receives synchronised streams of one-second mid-price returns, depth-of-book features, implied-volatility surfaces across moneyness 0.7–1.3, and realised-volatility estimators across rolling 5-minute and 30-minute windows. At market open the fused volatility encoder produces a latent state V_t^* with transitional-regime probability 0.62 and stressed-regime probability 0.28.

For an S&P 500 constituent trading at 412.5, the arbitrage-guarded ensemble generates a European call price of 7.84 for strike 415 and maturity 0.12 years, with delta 0.47 and calibrated uncertainty ± 0.21 . The Neural SDE block simulates 50,000 risk-neutral paths and prices a barrier contract with knock-out level $0.85 S_0$ at 4.31, with pathwise vega 0.19. The explainability layer attributes 41% of the price movement to short-term skew steepening and 27% to order-flow imbalance, and the constraint audit confirms no monotonicity violations and calendar-spread consistency within tolerance 10^{-4} .

Theory-guided calibration combines a spline-based local-volatility prior of 4.08 with a learned residual of 0.23 to produce a price interval [4.02, 4.56], with empirical coverage over recent rolling windows exceeding 95%. Using top-decile Precision@10 signals, the trading engine allocates across four low-correlation, regime-consistent stocks with momentum weights [0.31, 0.27, 0.22, 0.20]; over the following three hours, 146 trades execute at an average spread of 1.8 bps and turnover ratio 1.47, yielding gross/net returns of 0.67%/0.52%. One-day VaR is 1.36% with realised exceedance probability 4.1%, consistent with calibrated conformal bands, and delta-hedging error remains bounded at 0.061 during rolling rebalancing.

Under a subsequent volatility shock, directional accuracy falls from 78.6% to 74.2% with an eight-session regime-recovery period, limiting cumulative drawdown to 3.4% against 6.8% for policy-gradient baselines. This scenario illustrates how regime detection stabilises signals, volatility fusion improves price consistency, arbitrage enforcement protects surface integrity, dynamic path simulation enables coherent exotic valuation, and uncertainty calibration governs allocation and hedging decisions.

5 Conclusion and Future Scope

This paper proposed an integrated semi-parametric framework for regime-aware, arbitrage-consistent, and explainable high-frequency derivatives pricing. The five-module architecture — M³V-FuseNet, Arb-Guard, ExoX-SDE, A-XAI TraceStack, and TheoMix-CP — addresses five persistent gaps in the option-pricing literature: multi-source volatility fusion; near-hard no-arbitrage enforcement, exact on the continuum domain (Theorem 1) and approximated at finite grid resolution in the deployed system (Table 6, Corollary 1); surface-consistent exotic pricing; formal explainability with compliance diagnostics; and distribution-free uncertainty calibration.

Evaluation on 18 months of high-frequency equity-index and single-stock option data demonstrates European option RMSE reductions of 27–35% over Deep Hedging, IV surface RMSE of 1.28 bps, exotic RMSE reductions exceeding 35% for barrier/Asian/lookback payoffs, near-elimination of static arbitrage violations (0.06% residual, attributable to finite-grid discretisation rather than enforcement failure), and empirically observed conformal-prediction coverage of 95.8% — consistent with, though not identical to, the $\geq 95\%$ distribution-free guarantee, which holds under the exchangeability assumption discussed in Section 5.1. Ablation (Section 4.1) confirms statistically significant individual contributions from all five modules, with all improvements confirmed by Diebold–Mariano, Wilcoxon, paired t -tests, and Kupiec backtesting ($p < 0.005$). Future work will extend the framework to cross-asset macro conditioning, real-time news embeddings, and federated learning for privacy-preserving calibration across distributed derivatives desks.

5.1 Limitations

Despite strong empirical performance, several limitations apply. The architecture’s high-dimensional volatility fusion and relational state generation require rich intra-day data and continuous surface reconstruction, demanding substantial computing and storage resources; the system performs well for liquid markets but may not scale

to illiquid assets or fragmented markets with sparse microstructure data. Second, regime identification and probabilistic routing increase model complexity and risk over-parameterisation when training samples are limited or regime labels are poorly separable. Third, conformal calibration guarantees statistically correct coverage under exchangeability, which may be violated during prolonged crises or structural regulatory changes affecting small calibration windows. Fourth, transaction-cost modelling, while realistic, remains stylised and does not capture hidden liquidity, dark-pool execution dynamics, or adversarial market impact. Finally, the explainability layer provides structured diagnostics, but interpretation remains model-dependent and may not satisfy strict regulatory audit criteria in all jurisdictions.

5.2 Future Vision

The framework suggests several natural extensions. Incorporating cross-asset macro factors and real-time news embeddings into the volatility-fusion layer could improve regime anticipation and early shock detection during geopolitical or macroeconomic transitions. Multi-layer hierarchical graphs modelling sectoral, geographic, and derivative-underlying dependencies could support unified equity-derivatives-futures trading strategies. Adaptive execution policies that jointly optimise timing, sizing, and venue selection could reduce market impact and improve post-cost profitability. Embedding coherent risk measures such as conditional expected shortfall directly into the allocation layer could strengthen tail protection in extreme scenarios, while extending conformal calibration to time-adaptive, distribution-shift-aware variants could stabilise coverage during crises. Quantum-inspired optimisation and GPU-FPGA co-design offer a path to reduced latency for ultra-low-latency deployment. Finally, a long-horizon analysis across economic cycles, asset classes, and regulatory regimes would help establish the structural stability and governance readiness of the framework as a platform for next-generation autonomous financial systems that integrate learning, control, optimisation, and risk management in a transparent and auditable manner.

Acknowledgements. The authors would like to thank the Department of Information Technology and the Indian Institute of Information Technology, Nagpur, for providing access to computational infrastructure used in this research. The authors also acknowledge the data provider for access to the high-frequency derivatives dataset used in the experimental evaluation. The authors thank the anonymous reviewers whose constructive comments substantially improved the manuscript.

Declarations

- **Funding:** Not applicable.
- **Conflict of interest/Competing interests:** The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.
- **Ethics approval and consent to participate:** Not applicable.
- **Consent for publication:** Not applicable.

- **Data availability:** The raw high-frequency market data used in this study were obtained from a major electronic derivatives exchange under a commercial data licence agreement and cannot be publicly redistributed. The derived datasets generated during this study — including engineered volatility-state features, realised-volatility estimates, regime-classification labels, implied-volatility surface reconstructions, and model evaluation outputs — together with the complete model implementation code (M³V-FuseNet, Arb-Guard, ExoX-SDE, A-XAI TraceStack, TheoMix-CP), training scripts, hyperparameter configurations, and evaluation pipelines, are publicly available at: [GitHub Repository URL] and [Zenodo DOI: 10.5281/zenodo.XXXXXXX]. Researchers wishing to replicate the full pipeline with the original raw tick data may apply directly to the data provider.
- **Materials availability:** Not applicable.
- **Code availability:** See Data availability above.
- **Author contribution:** Chandrayani Rokde: Conceptualisation, Methodology, Software, Formal Analysis, Writing — Original Draft, Writing — Review & Editing, Visualisation. Jagdish Chakole: Conceptualisation, Software, Validation, Data Curation, Writing — Original Draft, Writing — Review & Editing, Supervision, Project Administration. Aishwarya Ukey: Writing — Review & Editing, Supervision, Project Administration.

References

- [1] Chen, R.T.Q., Li, B., Grosse, R., Duvenaud, D.: Neural ordinary differential equations applied to stochastic volatility. arXiv preprint arXiv:2104.03440 (2021) <https://doi.org/10.48550/arXiv.2104.03440>
- [2] Kidger, P., Foster, J., Li, X., Lyons, T.J.: Neural sdes as infinite-dimensional gans. In: Proceedings of the 38th International Conference on Machine Learning (ICML), vol. 139, pp. 5453–5463 (2021). <https://doi.org/10.48550/arXiv.2102.03657>
- [3] Buehler, H., Gonon, L., Teichmann, J., Wood, B.: Deep hedging. Quantitative Finance **19**(8), 1271–1291 (2019) <https://doi.org/10.1080/14697688.2019.1571683>
- [4] Black, F., Scholes, M.: The pricing of options and corporate liabilities. Journal of Political Economy **81**(3), 637–654 (1973) <https://doi.org/10.1086/260062>
- [5] Merton, R.C.: Option pricing when underlying stock returns are discontinuous. Journal of Financial Economics **3**, 125–144 (1976) [https://doi.org/10.1016/0304-405X\(76\)90022-2](https://doi.org/10.1016/0304-405X(76)90022-2)
- [6] Heston, S.L.: A closed-form solution for options with stochastic volatility. Review of Financial Studies **6**(2), 327–343 (1993) <https://doi.org/10.1093/rfs/6.2.327>
- [7] Bates, D.S.: Jumps and stochastic volatility. Review of Financial Studies **9**(1),

- 69–107 (1996) <https://doi.org/10.1093/rfs/9.1.69>
- [8] Hagan, P., Kumar, D., Lesniewski, A., Woodward, D.: Managing smile risk. *Wilmott Magazine*, 84–108 (2002). No DOI – Wilmott magazine article
 - [9] Bergomi, L.: Smile dynamics ii. *Risk*, 67–73 (2005) <https://doi.org/10.2139/ssrn.1493302>
 - [10] Dupire, B.: Pricing with a smile. *Risk* **7**(1), 18–20 (1994). No DOI – pre-DOI *Risk* magazine article
 - [11] Gatheral, J., Jacquier, A.: Arbitrage-free svi volatility surfaces. *Quantitative Finance* **14**(1), 59–71 (2014) <https://doi.org/10.1080/14697688.2013.819986>
 - [12] Cont, R., Fonseca, J.: Dynamics of implied volatility surfaces. *Quantitative Finance* **2**(1), 45–60 (2002) <https://doi.org/10.1088/1469-7688/2/1/304>
 - [13] Gatheral, J., Jaisson, T., Rosenbaum, M.: Volatility is rough. *Quantitative Finance* **18**(6), 933–949 (2018) <https://doi.org/10.1080/14697688.2017.1393551>
 - [14] Hernandez, A.: Model calibration with neural networks. *Risk*, 1–6 (2017) <https://doi.org/10.2139/ssrn.2812140> . DOI (SSRN preprint)
 - [15] Liu, S., Borovkykh, A., Grzelak, L.A., Oosterlee, C.W.: A neural network-based framework for financial model calibration. *Journal of Mathematics in Industry* **9**(9) (2019) <https://doi.org/10.1186/s13362-019-0066-7>
 - [16] Ackerer, D., Tagasovska, N., Vatter, T.: Deep smoothing of the implied volatility surface. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2020). <https://doi.org/10.48550/arXiv.1906.05065>
 - [17] Horvath, B., Muguruza, A., Tomas, M.: Deep learning volatility. *Quantitative Finance* **21**(1), 11–27 (2021) <https://doi.org/10.1080/14697688.2020.1817974>
 - [18] Antonov, A., Issakov, S., McClelland, A.: Efficient simm-mva calculations. *Risk Magazine* (2018) <https://doi.org/10.2139/ssrn.3040061> . DOI (SSRN)
 - [19] Ruf, J., Wang, W.: Neural networks for option pricing and hedging: A literature review. *Journal of Computational Finance* **24**(1), 1–46 (2020) <https://doi.org/10.21314/JCF.2020.390>
 - [20] Gierjatowicz, P., *et al.*: Robust pricing and hedging via neural stochastic differential equations. *Journal of Computational Finance* **26**(3), 1–32 (2022) <https://doi.org/10.21314/JCF.2022.025>
 - [21] Cuchiero, C., Khosrawi, W., Teichmann, J.: A generative adversarial network approach to calibration of local stochastic volatility models. *Risks* **8**(4), 101 (2020) <https://doi.org/10.3390/risks8040101>

- [22] Wiese, M., Knobloch, R., Korn, R., Kretschmer, P.: Quant gans. *Quantitative Finance* **20**(9), 1419–1440 (2020) <https://doi.org/10.1080/14697688.2020.1730426>
- [23] Cao, J., Chen, J., Hull, J., Poulos, Z.: Ai and machine learning in accounting and finance. *Journal of Accounting Literature* (2021). No verifiable DOI located
- [24] De Spiegeleer, J., Madan, D.B., Reyners, S., Schoutens, W.: Machine learning for quantitative finance. *Quantitative Finance* **18**(10), 1635–1643 (2018) <https://doi.org/10.1080/14697688.2018.1495335>
- [25] Angelopoulos, A.N., Bates, S.: A gentle introduction to conformal prediction. *arXiv preprint arXiv:2107.07511* (2023) <https://doi.org/10.48550/arXiv.2107.07511>
- [26] Chen, R.T.Q., Rubanova, Y., Bettencourt, J., Duvenaud, D.: Neural ordinary differential equations. In: *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 31 (2018). <https://doi.org/10.48550/arXiv.1806.07366>
- [27] Li, X., Wong, T.-K.L., Chen, R.T.Q., Duvenaud, D.: Scalable gradients for stochastic differential equations. In: *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS)*, vol. 108, pp. 3870–3882 (2020). <https://doi.org/10.48550/arXiv.2001.01328>
- [28] Diebold, F.X., Mariano, R.S.: Comparing predictive accuracy. *Journal of Business & Economic Statistics* **13**(3), 253–263 (1995) <https://doi.org/10.1080/07350015.1995.10524599>
- [29] Kupiec, P.H.: Techniques for verifying the accuracy of risk measurement models. *Journal of Derivatives* **3**(2), 73–84 (1995) <https://doi.org/10.3905/jod.1995.407942>
- [30] Carr, P., Ellis, K., Gupta, V.: Static hedging of exotic options. *Journal of Finance* **53**(3), 1165–1190 (1998) <https://doi.org/10.1111/0022-1082.00048>
- [31] Derman, E., Kani, I.: Riding on a smile. *Risk* **7**(2), 32–39 (1994). No DOI – pre-DOI Risk magazine article
- [32] Andersen, L., Brotherton-Ratcliffe, R.: The equity option volatility smile. *Journal of Computational Finance* **1**(2) (1998). No DOI located – pre-DOI era journal issue
- [33] Bayer, C., Friz, P., Gatheral, J.: Pricing under rough volatility. *Quantitative Finance* **16**(6), 887–904 (2016) <https://doi.org/10.1080/14697688.2015.1099717>
- [34] Salvi, C., *et al.*: Higher order kernel mean embeddings. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2021). <https://doi.org/10.48550/arXiv.2109.03582>

- [35] Rudin, W.: Real and Complex Analysis, 3rd edn. McGraw-Hill, New York (1987).
No DOI – pre-DOI classic textbook
- [36] Nocedal, J., Wright, S.J.: Numerical Optimization, 2nd edn. Springer, New York (2006). <https://doi.org/10.1007/978-0-387-40065-5>

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [supplementarymaterials.docx](#)